

# Real-time semiparametric regression

BY J. LUTS<sup>1</sup>, T. BRODERICK<sup>2</sup> AND M.P. WAND<sup>1</sup>

<sup>1</sup> School of Mathematical Sciences, University of Technology Sydney, Broadway 2007, Australia

<sup>2</sup> Department of Statistics, University of California, Berkeley, California 94720, USA

17th September, 2012

## SUMMARY

We develop algorithms for performing semiparametric regression analysis in real time, with data processed as it is collected and made immediately available via modern telecommunications technologies. Our definition of semiparametric regression is quite broad and includes, as special cases, generalized linear mixed models, generalized additive models, geostatistical models, wavelet nonparametric regression models and their various combinations. Fast updating of regression fits is achieved by couching semiparametric regression into a Bayesian hierarchical model or, equivalently, graphical model framework and employing online mean field variational ideas. Internet sites attached to this article illustrate the methodology for continually arriving stock market and real estate data. Flexible real-time analyses, based on increasingly ubiquitous streaming data sources stand to benefit.

*Keywords:* Approximate Bayesian inference; Generalized additive models; Mean field variational Bayes; Mixed models; Online variational Bayes; Penalized splines; Wavelets.

## 1 Introduction

Ongoing technological advancements mean that data are being collected and made available for inference with rapidly increasing volume and speed. There are numerous examples of this explosion of data, but two that have established connections with semiparametric regression, our focus in this article, are Internet auction analysis (e.g. Jank & Shmueli, 2007) and real-time spatial epidemiology (e.g. Kaimi & Diggle, 2011).

*Semiparametric regression* refers to a large class of regression models that provide for non-linear predictor effects using spline and wavelet basis functions, as well as dependencies arising in grouped data such as within-subject correlation. An arsenal of both frequentist and Bayesian fitting and inference procedures now exist. Recent overviews are contained in Ruppert, Wand & Carroll (2009) and Wand & Ormerod (2011).

Virtually all semiparametric regression methodology proposed to date assume that the data are processed in *batch*. That is, each iteration of an iterative inference procedure requires statistics computed across the entire data set. Summaries such as function estimates, confidence intervals and posterior density functions are then outputted. Downsides to batch processing include the requirement that statistical analysis wait until an entire data set has been assembled and, sometimes, the necessity of storing the entire data set in memory. In the *online* case, each iteration of the procedure requires only knowledge of the latest data point and any summary statistics from previous iterations. A particular advantage of online processing is that summaries, such as those just mentioned, are updated throughout the data collection process and therefore are available immediately upon demand. Online processing also has the advantage of not requiring storage of potentially very large data-sets.

While a number of batch procedures exist for performing semiparametric regression, we focus on a particular methodology here due to the ease of adapting it to the online

framework as well as its wide range of applicability. Consider single predictor nonparametric regression, a special case of semiparametric regression with a long history and large literature. Fully automatic nonparametric regression batch procedures include: (a) local linear kernel smoother with cross-validation bandwidth selection, (b) local linear kernel smoother with direct plug-in bandwidth selection, (c) frequentist low-rank smoothing spline with restricted maximum likelihood smoothing parameter selection, (d) Bayesian low-rank smoothing spline with Markov chain Monte Carlo approximate inference and (e) Bayesian low-rank smoothing spline with mean field variational Bayesian (MFVB) approximate inference. Details of (a) are in Härdle (1990), details of (b) are in Wand & Jones (1995), whilst (c) and (d) are described in Ruppert, Wand & Carroll (2003). Section 2.7 of Wand & Ormerod (2011) explains (e). Approaches (a)–(d) are more established, but none have a viable online modification. However, (e) is relatively easy to modify for this purpose.

Another advantage of the Bayesian low-rank smoothing spline approach to nonparametric regression is its extendibility. As explained in Wand (2009), couching semiparametric regression in a graphical models framework permits arbitrarily sophisticated models to be handled elegantly, efficiently, and cohesively. This approach can handle generalized additive models, geostatistical models, wavelet nonparametric regression models and their various combinations, as well complications such as outliers and missingness. Inference in these models is often accomplished by applying Markov chain Monte Carlo procedures using the directed acyclic graph of variable dependencies. While versatile and accurate, such inference procedures can be unacceptably slow. MFVB approaches, as demonstrated in Faes, Ormerod & Wand (2011) and Wand & Ormerod (2011), are a much faster alternative. Some accuracy and versatility must be sacrificed in return for the increased speed of MFVB. Nonetheless, for the models treated in this article MFVB accuracy ranges from good to excellent.

Iterative algorithms that make a single pass through the data, with one iteration per data point, have recently been developed for variational Bayesian inference. In the machine learning literature, Hoffman, Blei & Bach (2010) introduced such an MFVB algorithm for latent Dirichlet allocation and applied their algorithm to topic modeling. This procedure was extended to the hierarchical Dirichlet process by Wang, Paisley & Blei (2011). Tchumtchoua, Dunson & Morris (2012) further developed online MFVB approximate inference for high-dimensional correlated data. These articles are referred to as *online mean field variational Bayes* or often with the shorter name *online variational Bayes*. While they are indeed single-pass and require storing only one data point in memory, they do, however, require knowledge of the number of data points from the start of the algorithm. Our focus in this work, by contrast, is not on transforming MFVB algorithms that require multiple data passes into single-pass algorithms. Rather, we are, in some sense, pursuing a more classical definition of an “online algorithm” in that each iteration of our procedure uses past data only in the form of sufficient statistics and future data not at all.

Online MFVB has not been entertained previously for nonparametric and semiparametric regression. But there is an old and large literature involving other online approaches. For the nonparametric regression and the related density estimation problem Wolverton & Wagner (1967), Yamato (1971), Devroye & Wagner (1980) and Krzyak & Pawlak (1984) are examples of early articles on online analysis using kernel estimators. However, they are chiefly concerned with theoretical properties of the estimators and are devoid of practical automatic smoothing parameter selection strategies.

Outside of semiparametric regression there are also large literatures on online analysis. A few recent examples are: Ng, McLachlan & Lee (2006) on prediction of hospital resource utilization, Fricker & Chang (2008) on biosurveillance and Kaimi & Diggle (2011) on monitoring of variation in risk of infections. A very recent article by Michalak *et al.* (2012) describes the development of systems for real-time streaming analysis.

Semiparametric regression is a highly visual branch of Statistics, with graphics being

a crucial means of conveying and diagnosing regression fits. The norm for such graphical display are ink drawings on pieces of paper or figures in PDF file. Real-time semiparametric regression represents a paradigm shift in graphical display, where regression summaries are best thought of as dynamic graphics on web-pages or iDevice apps. We have organized an Internet site that illustrates real-time semiparametric regression graphical display.

Section 2 introduces the notion of real-time semiparametric regression based online MFVB via increasingly more sophisticated Gaussian response models. Both classical and sparse shrinkage are treated. The more challenging binary response case is dealt with in Section 3. Dynamic web-pages that illustrate the new methodology on live data are the focus of Section 4.

## 2 Gaussian Response Models

The conversion of a batch MFVB semiparametric regression procedure to one that does online processing is particularly straightforward in the Gaussian response case. We start by explaining such conversion for the multiple linear regression model, since it has minimal notational overhead.

### 2.1 Multiple linear regression

Let  $\mathbf{X}$  be a  $p \times n$  design matrix and consider the Bayesian regression model

$$\mathbf{y} | \boldsymbol{\beta}, \sigma \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}), \quad \boldsymbol{\beta} \sim N(\mathbf{0}, \sigma_{\beta}^2 \mathbf{I}), \quad \sigma \sim \text{Half-Cauchy}(A). \quad (1)$$

An equivalent, but more tractable model, is that where  $\sigma \sim \text{Half-Cauchy}(A)$  is replaced by the auxiliary variable representation

$$\sigma^2 | a \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a), \quad a \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/A^2). \quad (2)$$

Figure 1 displays the directed acyclic graph corresponding to the model conveyed by (1) and (2).

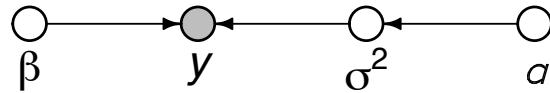


Figure 1: *Directed acyclic graph for the model conveyed by (1) and (2). The shading corresponds to the observed data.*

If the joint posterior density function  $p(\boldsymbol{\beta}, \sigma^2, a | \mathbf{y})$  is restricted to having the product form

$$q(\boldsymbol{\beta}, a) q(\sigma^2) \quad (3)$$

for density functions  $q(\boldsymbol{\beta}, a)$  and  $q(\sigma^2)$  then the optimal density functions in terms of minimum Kullback-Leibler distance can be found from the batch MFVB algorithm, Algorithm 1. The optimal density functions are denoted by  $q^*(\boldsymbol{\beta}, a)$  and  $q^*(\sigma^2)$ . The MFVB solution is also such that  $q^*(\boldsymbol{\beta}, a) = q^*(\boldsymbol{\beta}) q^*(a)$  even though (3) does not assume this. The approximate marginal likelihood induced by product restriction (3) is denoted by  $\underline{p}(\mathbf{y}; q)$  and is a lower bound on the exact marginal likelihood  $p(\mathbf{y})$ . See, for example, Section 2.2 of Ormerod & Wand (2010) for the underpinning of Algorithm 1. General references on MFVB include Bishop (2006) and Wainwright & Jordan (2008).

---

**Algorithm 1** *Batch mean field variational Bayes algorithm for approximate inference in the Gaussian response linear regression model (1) and (2).*

---

Initialize:  $\mu_{q(1/\sigma^2)} > 0$ .

Read in  $\mathbf{y}$  ( $n \times 1$ ) and  $\mathbf{X}$  ( $n \times p$ ).

Cycle:

$$\begin{aligned}\boldsymbol{\Sigma}_{q(\beta)} &\leftarrow \left\{ \mu_{q(1/\sigma^2)} \mathbf{X}^T \mathbf{X} + \sigma_{\beta}^{-2} \mathbf{I} \right\}^{-1} \\ \boldsymbol{\mu}_{q(\beta)} &\leftarrow \mu_{q(1/\sigma^2)} \boldsymbol{\Sigma}_{q(\beta)} \mathbf{X}^T \mathbf{y} ; \mu_{q(1/a)} \leftarrow 1 / \{\mu_{q(1/\sigma^2)} + A^{-2}\} \\ \mu_{q(1/\sigma^2)} &\leftarrow \frac{n+1}{2\mu_{q(1/a)} + \mathbf{y}^T \mathbf{y} - 2\boldsymbol{\mu}_{q(\beta)}^T \mathbf{X}^T \mathbf{y} + \text{tr}[(\mathbf{X}^T \mathbf{X}) \{\boldsymbol{\Sigma}_{q(\beta)} + \boldsymbol{\mu}_{q(\beta)} \boldsymbol{\mu}_{q(\beta)}^T\}]}\end{aligned}$$

until the increase in  $\underline{p}(\mathbf{y}; q)$  is negligible.

Produce summaries based on  $q^*(\beta) \sim N(\boldsymbol{\mu}_{q(\beta)}, \boldsymbol{\Sigma}_{q(\beta)})$  and

$$q^*(\sigma^2) \sim \text{Inverse-Gamma}(\frac{1}{2}(n+1), (n+1)/\{2\mu_{q(1/\sigma^2)}\}).$$


---

The approximate marginal log-likelihood, used to monitor convergence in Algorithm 1, has explicit expression

$$\begin{aligned}\log \underline{p}(\mathbf{y}; q) &= \frac{1}{2} p - \frac{1}{2} n \log(2\pi) - 2 \log(\pi) + \log \Gamma(\frac{1}{2}(n+1)) \\ &\quad - \frac{1}{2} p \log(\sigma_{\beta}^2) - \log(A) - \frac{1}{2\sigma_{\beta}^2} \{\|\boldsymbol{\mu}_{q(\beta)}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\beta)})\} \\ &\quad + \frac{1}{2} \log |\boldsymbol{\Sigma}_{q(\beta)}| - \frac{1}{2}(n+1) \log[(n+1)/\{2\mu_{q(1/\sigma^2)}\}] \\ &\quad - \log(\mu_{q(1/\sigma^2)} + A^{-2}) + \mu_{q(1/\sigma^2)} \mu_{q(1/a)}.\end{aligned}$$

In Algorithm 1, dependence on the data is only through the quantities  $\mathbf{y}^T \mathbf{y}$ ,  $\mathbf{X}^T \mathbf{y}$  and  $\mathbf{X}^T \mathbf{X}$  and each of these have simple updates when a new response  $y_{\text{new}}$  and its corresponding  $p \times 1$  vector of predictors  $\mathbf{x}_{\text{new}}$  arrives. For example, the new  $\mathbf{X}^T \mathbf{X}$  matrix is

$$\mathbf{X}_{\text{new}}^T \mathbf{X}_{\text{new}} = \mathbf{X}^T \mathbf{X} + \mathbf{x}_{\text{new}} \mathbf{x}_{\text{new}}^T.$$

Based on these observations Algorithm 2, the *online* modification of the Algorithm 1, ensues.

Algorithm 2 differs from Algorithm 1 in that the data are processed on arrival and the approximate posterior densities of the model parameters are continually updated. In the case of streaming data there is the option of dynamic graphical displays of the approximate posterior density functions of the regression coefficients and error variance and corresponding approximate Bayes estimates and credible sets. Dynamic regression diagnostic plots could also be entertained.

Figure 2 provides rudimentary illustration of online regression inference when data from the Vietnam World Bank Living Standards Survey (source: Cameron & Trivedi, 2005) is fed into Algorithm 2. These data are in the `VietNamI` data-frame of the R package `Ecdat` (Croissant, 2011). The response variable is the logarithm of total medical expenses. Description of the predictor variables is given in the `VietNamI` documentation of Croissant (2011). Each variable was transformed to lie inside the unit interval before being processed. Note that the scaling was determined using an initialization batch just as for the initial parameter tuning; therefore, the single-pass aspect of the algorithm is maintained. The posterior density functions were then back-transformed to correspond to the

---

**Algorithm 2** *Online mean field variational Bayes algorithm for approximate inference in the Gaussian response linear regression model (1) and (2).*

---

Initialize:  $\mu_{q(1/\sigma^2)} > 0$ ,  $\mathbf{y}^T \mathbf{y} \leftarrow 0$ ,  $\mathbf{X}^T \mathbf{y} \leftarrow \mathbf{0}_{p \times 1}$ ,  $\mathbf{X}^T \mathbf{X} \leftarrow \mathbf{0}_{p \times p}$ ,  $n \leftarrow 0$ .

Cycle:

read in  $y_{\text{new}}$  ( $1 \times 1$ ) and  $\mathbf{x}_{\text{new}}$  ( $p \times 1$ ) ;  $n \leftarrow n + 1$   
 $\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}^T \mathbf{y} + y_{\text{new}}^2$  ;  $\mathbf{X}^T \mathbf{y} \leftarrow \mathbf{X}^T \mathbf{y} + \mathbf{x}_{\text{new}} y_{\text{new}}$  ;  $\mathbf{X}^T \mathbf{X} \leftarrow \mathbf{X}^T \mathbf{X} + \mathbf{x}_{\text{new}} \mathbf{x}_{\text{new}}^T$   
 $\Sigma_{q(\beta)} \leftarrow \left\{ \mu_{q(1/\sigma^2)} \mathbf{X}^T \mathbf{X} + \sigma_{\beta}^{-2} \mathbf{I} \right\}^{-1}$   
 $\mu_{q(\beta)} \leftarrow \mu_{q(1/\sigma^2)} \Sigma_{q(\beta)} \mathbf{X}^T \mathbf{y}$  ;  $\mu_{q(1/a)} \leftarrow 1 / \{\mu_{q(1/\sigma^2)} + A^{-2}\}$   
 $\mu_{q(1/\sigma^2)} \leftarrow \frac{n + 1}{2\mu_{q(1/a)} + \mathbf{y}^T \mathbf{y} - 2\mu_{q(\beta)}^T \mathbf{X}^T \mathbf{y} + \text{tr}[(\mathbf{X}^T \mathbf{X})\{\Sigma_{q(\beta)} + \mu_{q(\beta)} \mu_{q(\beta)}^T\}]}$   
produce summaries based on  $q^*(\beta) \sim N(\mu_{q(\beta)}, \Sigma_{q(\beta)})$  and  
 $q^*(\sigma^2) \sim \text{Inverse-Gamma}(\frac{1}{2}(n + 1), (n + 1) / \{2\mu_{q(1/\sigma^2)}\})$

until data no longer available or analysis terminated.

---



---

original units. The hyperparameters were set at  $\sigma_{\beta}^2 = 10^{10}$  and  $A = 10^5$  to impose non-informativity.

Note, for example, the approximate posterior density functions for  $\beta_2$ , the regression coefficient attached to age of household head. For  $n \leq 100$  the posterior density function is relatively flat and  $\beta_2$  is not statistically significant. As  $n$  increases, the posterior density functions become narrower and, by  $n = 250$ , the lower limit of the 95% credible set is positive – indicating statistical significance of this predictor. Also, in the bottom row, the posterior mean of the error variance,  $\sigma^2$ , is seen to gradually move to the left as  $n$  increases – reflecting the increased precision as more data arrives.

The right-most column of Figure 2 shows the batch MFVB posterior density functions for  $n = 250$ . In this case, the batch and online MFVB results are seen to be virtually identical. However, as demonstrated later, such agreement is not guaranteed in general.

### 2.1.1 Batch-based Tuning and Convergence Diagnosis

Ideally, the online Algorithm 2 will mimic the results of the batch Algorithm 1 as the sample size  $n$  increases. However, we know of no guarantees that this will happen and it is possible that the online parameters will diverge from their batch counterparts. For the more elaborate models studied later in this article, such divergence is very common. Therefore, convergence diagnosis at the start of the online iterations is essential. The principal idea is to start by running a small subset of initial data points in the batch algorithm to obtain starting values for both data sufficient statistics and, more importantly, estimated parameters of the model. A second, small validation subset of data is used to compare the batch and online algorithm results. If convergence of the online iterations to their batch counterparts is not verified by this comparison then larger initial batch runs are required to tune the online algorithm.

The idea of collecting some of the data into a small subset to improve performance of an algorithm where one new data point is introduced at each iteration is reminiscent of the “mini-batches” of Hoffman, Blei & Bach (2010). However, in our approach, the batching of data happens only with a small subset at the very beginning of the algorithm rather than throughout. Also, we develop an alternative tuning method for this subset batch size

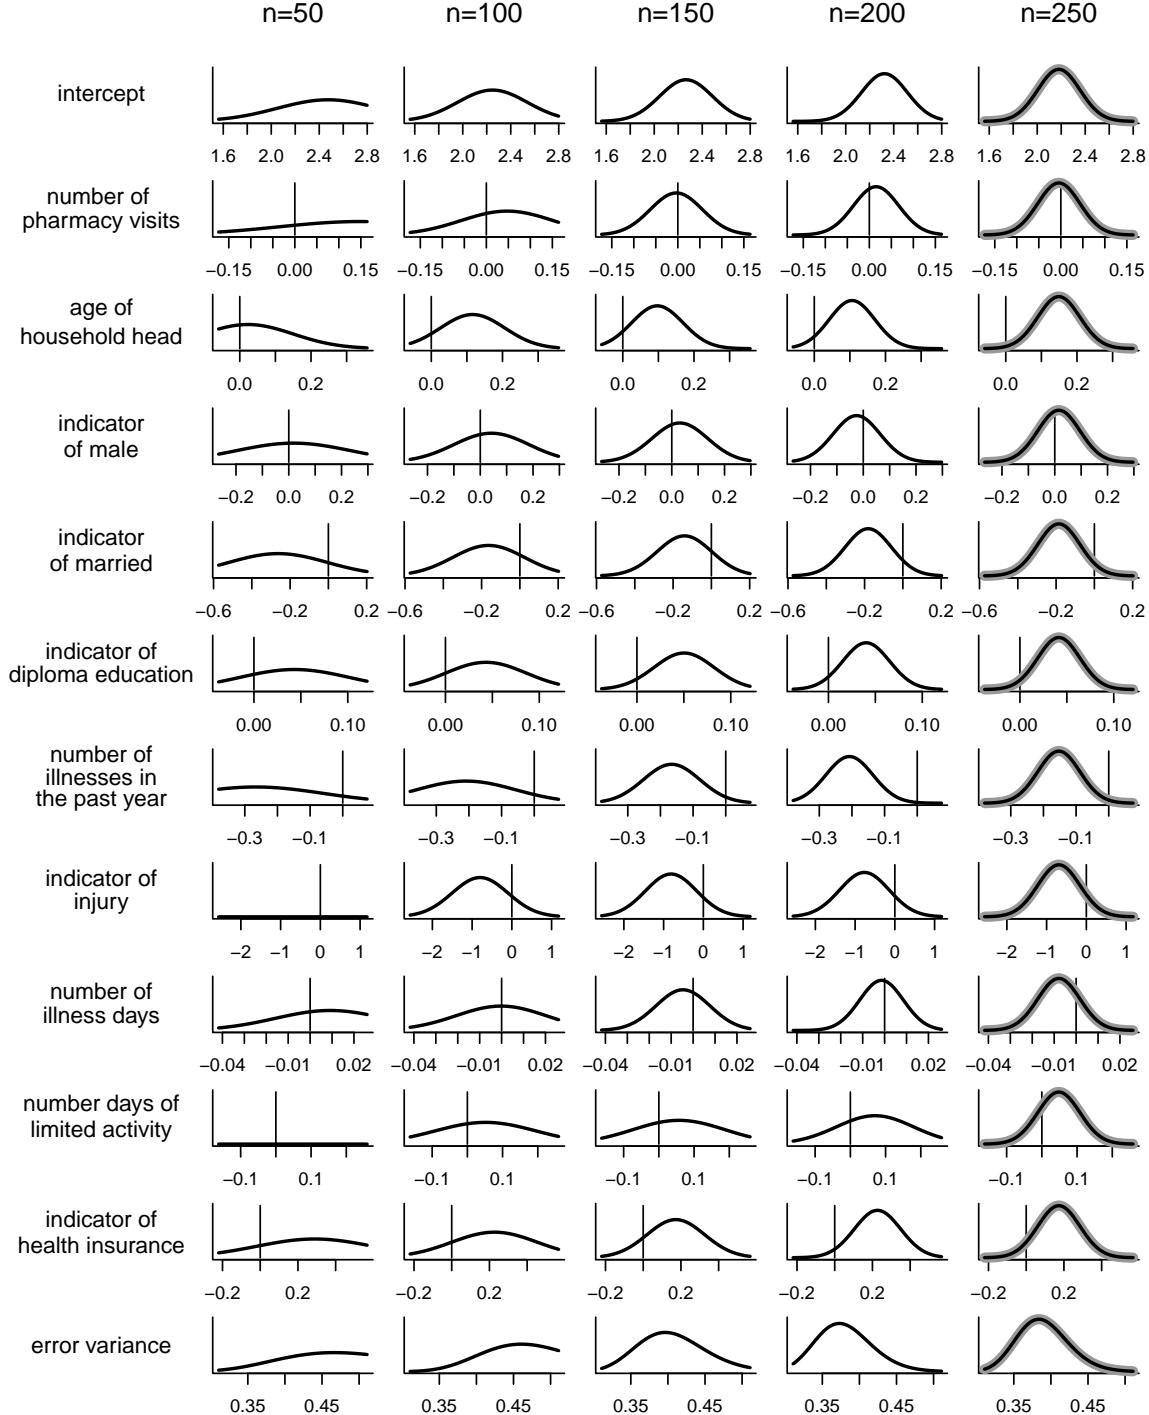


Figure 2: Successive approximate posterior density functions of regression coefficients and the logarithm of error variance for the Vietnam medical expenses data described in the text. The predictors corresponding to each regression coefficient are listed in the left-hand columns. The posterior density functions are based on online MFVB as detailed in Algorithm 2. The axis limits are the same across each row and a vertical line is positioned at zero. For  $n = 250$  the batch MFVB approximate fits are shown as thick grey curves.

below; notably, our tuning method requires batching only some initial subset of the data rather than the full data set.

We will now provide details via the Figure 2 example. Figure 3 shows the posterior means and 95% credible sets for each  $\beta_j$ ,  $0 \leq j \leq 11$ , and  $\log(\sigma^2)$  and samples sizes

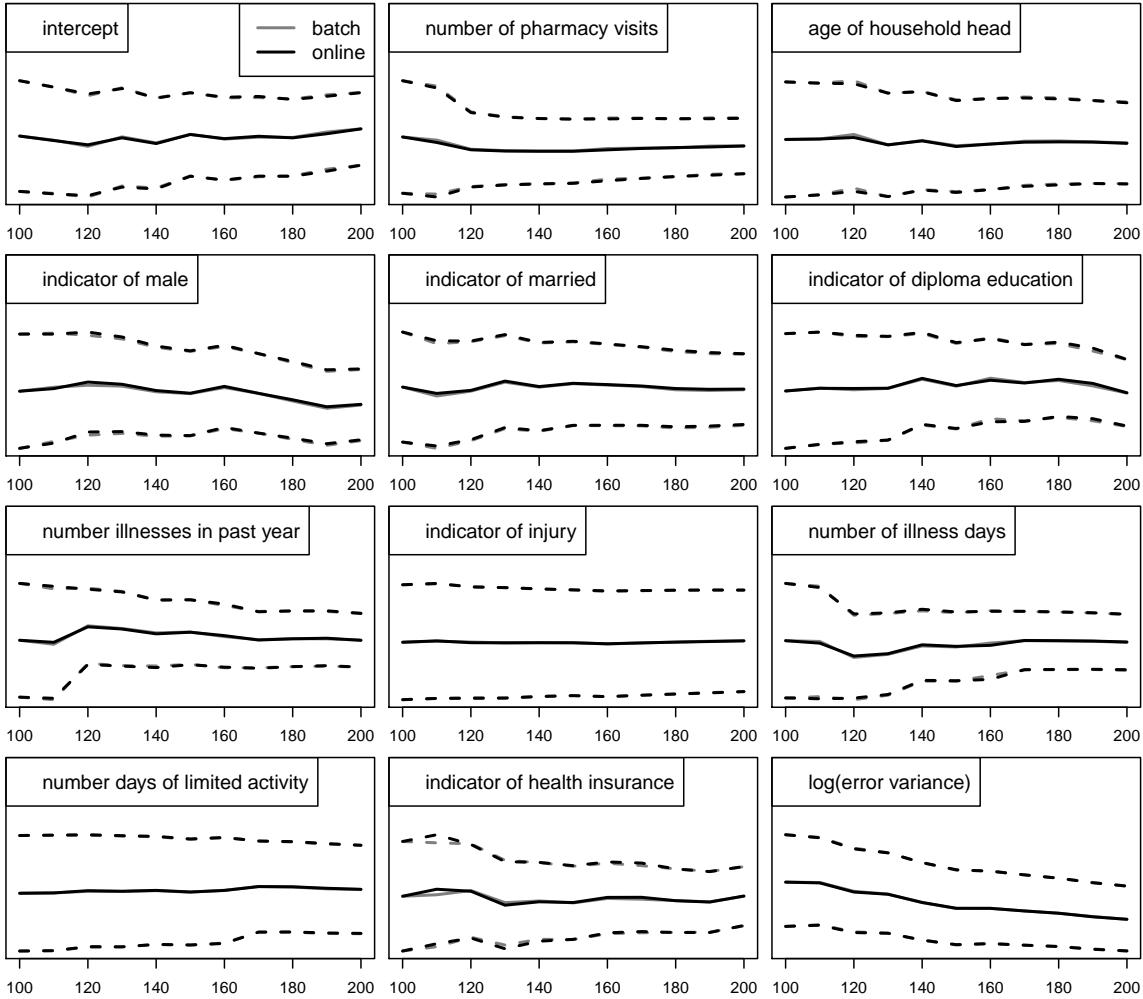


Figure 3: Convergence diagnostics for the example given in Figure 2. The horizontal axes show the sample sizes between a warm-up batch sample of size  $n_{\text{warm}} = 100$  and validation sample sizes up to  $n_{\text{valid}} = 100$  greater than  $n_{\text{warm}}$ .

$n = 100, 110, \dots, 200$  when the Vietnam medical expenses data are fitted via both batch and online MFVB. The batch MFVB summary statistics (shown as grey lines in Figure 3) correspond to simply inputting the first  $n_{\text{warm}} = 100$  observations into Algorithm 1 and then repeating this process for 10 additional equally-spaced sample sizes that are  $n_{\text{valid}} = 100$  greater than  $n_{\text{warm}}$ . The largest sample size is then  $n_{\text{warm}} + n_{\text{valid}} = 200$ . The online results (shown as grey lines in Figure 3) were obtained via online MFVB updating steps of Algorithm 2 but with  $\mu_{q(1/\sigma^2)}$ ,  $\mathbf{y}^T \mathbf{y}$ ,  $\mathbf{X}^T \mathbf{y}$ ,  $\mathbf{X}^T \mathbf{X}$  and  $n$  initialized at the values obtained when the first  $n_{\text{warm}} = 100$  observations are inputted into Algorithm 1. This implies that all the results are identical at  $n = n_{\text{warm}} = 100$ , but there are some small discrepancies for  $n > 100$ . In this example the discrepancies are negligible, and hard to discern from Figure 3 – indicating convergence of the online MFVB algorithm. Figure 5 in Section 3 shows an example where convergence is not achieved with  $n_{\text{warm}} = 100$  and a larger warm-up is required.

Algorithm 2' is a modification of Algorithm 2 that incorporates batch-based tuning and convergence diagnostics. Whilst such modification it is not necessary for the example depicted in Figures 2 and 3, it is crucial for more sophisticated semiparametric models such as those described later in this article.

---

1. Set  $n_{\text{warm}}$  to be the warm-up sample size and  $n_{\text{valid}}$  to be size of the validation period. Read in the first  $n_{\text{warm}} + n_{\text{valid}}$  response and predictor values.
2. Create  $\mathbf{y}_{\text{warm}}$  and  $\mathbf{X}_{\text{warm}}$  consisting of the first  $n_{\text{warm}}$  response and predictor values.
3. Feed  $\mathbf{y}_{\text{warm}}$  and  $\mathbf{X}_{\text{warm}}$  into the batch MFVB Algorithm 1 to obtain a starting value for  $\mu_{q(1/\sigma^2)}$ .
4. Set  $\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{X}^T \mathbf{y} \leftarrow \mathbf{X}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{X}^T \mathbf{X} \leftarrow \mathbf{X}_{\text{warm}}^T \mathbf{X}_{\text{warm}}$  and  $n \leftarrow n_{\text{warm}}$ .
5. Run the online MFVB Algorithm 2 until  $n = n_{\text{warm}} + n_{\text{valid}}$ .
6. Use convergence diagnostic graphics to assess whether the online parameters are converging to the batch parameters.
  - (a) If not converging then return to Step 1 and increase  $n_{\text{warm}}$ .
  - (b) If converging then continue running the online MFVB Algorithm 2 until data no longer available or analysis terminated.

---

Algorithm 2': *Modification of Algorithm 2 to include batch-based tuning and convergence diagnosis.*

One could contemplate automating Step 6 of Algorithm 2', to save the user from having to conduct diagnostic checks. However, we have not yet explored automatic convergence diagnosis and, instead, flag this as a problem worthy of future research.

### 2.1.2 Model assumptions

The online MFVB Algorithms 2' is founded upon the same assumptions as its batch counterpart Algorithm 1. Both algorithms fit the Bayesian linear regression model (1), but the latter has the option to do the fitting in real time for sequentially arriving data.

Throughout this article, we are not allowing for the model parameters to change as new data arrive. Colloquially, we assume ‘fixed targets’ rather than ‘moving targets’. Extensions to semiparametric regression scenarios where the model parameters drift over time, and real-time algorithms that adapt to such drifts, is certainly worthy of future investigation – but beyond this article’s scope.

## 2.2 Linear Mixed Models

A very useful structure for semiparametric regression is the class of Bayesian linear mixed models of the form

$$\begin{aligned} \mathbf{y} | \boldsymbol{\beta}, \mathbf{u}, \sigma_{\varepsilon}^2 &\sim N(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \sigma_{\varepsilon}^2 \mathbf{I}) \\ \mathbf{u} | \sigma_{u1}^2, \dots, \sigma_{ur}^2 &\sim N(\mathbf{0}, \text{blockdiag}(\sigma_{u1}^2 \mathbf{I}_{K_1}, \dots, \sigma_{ur}^2 \mathbf{I}_{K_r})) \end{aligned} \quad (4)$$

where  $\mathbf{y}$  is an  $n \times 1$  vector of response variables,  $\boldsymbol{\beta}$  is a  $p \times 1$  vector of fixed effects,  $\mathbf{u}$  is a vector of random effects,  $\mathbf{X}$  and  $\mathbf{Z}$  corresponding design matrices,  $\sigma_{\varepsilon}^2$  is the error variance and  $\sigma_{u1}^2, \dots, \sigma_{ur}^2$  are variance parameters corresponding to sub-blocks of  $\mathbf{u}$  of size  $K_1, \dots, K_r$ . Here the priors are taken to be

$$\boldsymbol{\beta} \sim N(\mathbf{0}, \sigma_{\beta}^2 \mathbf{I}), \quad \sigma_{u\ell} \sim \text{Half-Cauchy}(A_{u\ell}), \quad 1 \leq \ell \leq r, \quad \sigma_{\varepsilon} \sim \text{Half-Cauchy}(A_{\varepsilon}) \quad (5)$$

with the hyperparameters satisfying  $\sigma_\beta^2, A_\varepsilon, A_{u\ell} > 0$  for  $1 \leq \ell \leq r$ . As in Section 2, tractability considerations motivate the introduction of auxiliary variables

$$a_{u\ell} \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/A_{u\ell}^2) \quad \text{and} \quad a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/A_\varepsilon^2) \quad (6)$$

and use of the analogue of (2) to induce Half-Cauchy priors on the standard deviation parameters.

As spelt out in Section 2 of Zhao, Staudenmayer, Coull & Wand (2006), model (4)–(5) encompasses a rich class of models including (with example number from Zhao *et al.* 2006 added):

- simple random effects models (Examples 1 and 2),
- cross random effects models (Example 3),
- nested random effects models (Example 4),
- generalized additive models (Example 6),
- semiparametric mixed models (Example 7),
- bivariate smoothing and geoadditive models extensions (Example 8).

Examples 2 and 6 of Zhao *et al.* (2006) actually involve  $2 \times 2$  and  $3 \times 3$  unstructured covariance matrix parameters which are not covered by (5) and Algorithm 3. However, as discussed in Section 2.3, the unstructured covariance matrix extension is quite straightforward.

Algorithm 3 describes online MFVB fitting of the linear mixed model (4)–(6) with product restriction

$$q(\boldsymbol{\beta}, \mathbf{u}, a_{u1}, \dots, a_{ur}, a_\varepsilon, \sigma_{u1}^2, \dots, \sigma_{ur}^2, \sigma_\varepsilon^2) = q(\boldsymbol{\beta}, \mathbf{u}, a_{u1}, \dots, a_{ur}, a_\varepsilon) q(\sigma_{u1}^2, \dots, \sigma_{ur}^2, \sigma_\varepsilon^2).$$

Batch MFVB fitting of (4), but with slightly different prior distributions, is given by Algorithm 3 of Ormerod & Wand (2010), where the notation

$$\mathbf{C} = [\mathbf{X} \ \mathbf{Z}]$$

is used. Let  $P$  be the number of columns in  $\mathbf{C}$ . Then each pass of the corresponding online MFVB algorithm involves arrival and processing of a new scalar response measurement,  $y_{\text{new}}$ , and a  $P \times 1$ ,  $\mathbf{c}_{\text{new}}$ , corresponding to the new row of  $\mathbf{C}$ . Algorithm 3 facilitates online fitting of (4) via natural modification of the aforementioned batch MFVB algorithm.

The  $\mathbf{c}_{\text{new}}$  vector will have different forms depending on the type of linear mixed model. To better understand the nature these forms, consider the following two special cases of (4):

$$\begin{aligned} y_{ij} | \beta_0, U_i, \beta_1 &\stackrel{\text{ind.}}{\sim} N(\beta_0 + U_i + \beta_1 x_{ij}, \sigma_\varepsilon^2), \quad 1 \leq i \leq m, \quad 1 \leq j \leq n_i, \\ U_i | \sigma_u^2 &\stackrel{\text{ind.}}{\sim} N(0, \sigma_u^2 \mathbf{I}), \quad \beta_0, \beta_1 \stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2), \\ \sigma_u &\sim \text{Half-Cauchy}(A_u), \quad \sigma_\varepsilon \sim \text{Half-Cauchy}(A_\varepsilon) \end{aligned} \quad (7)$$

and

$$\begin{aligned} y_i | \beta_0, \beta_s, \beta_t, \mathbf{u}_s, \mathbf{u}_t &\stackrel{\text{ind.}}{\sim} N\left(\beta_0 + \beta_s s_i + \beta_t t_i + \sum_{k=1}^{K_s} u_{s,k} z_k^s(s_i) + \sum_{k=1}^{K_t} u_{t,k} z_k^t(t_i), \sigma_\varepsilon^2\right), \\ 1 \leq i \leq n, \quad \mathbf{u}_s &= [u_{s,1} \dots u_{s,K_s}]^T, \quad \mathbf{u}_t = [u_{t,1} \dots u_{t,K_t}]^T, \\ \mathbf{u}_s | \sigma_{u,s}^2 &\sim N(0, \sigma_{u,s}^2 \mathbf{I}), \quad \mathbf{u}_t | \sigma_{u,t}^2 \sim N(0, \sigma_{u,t}^2 \mathbf{I}), \\ \beta_0, \beta_s, \beta_t &\stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2), \quad \sigma_{u,s} \sim \text{Half-Cauchy}(A_{u,s}), \quad \sigma_{u,t} \sim \text{Half-Cauchy}(A_{u,t}), \\ \sigma_\varepsilon &\sim \text{Half-Cauchy}(A_\varepsilon). \end{aligned} \quad (8)$$

---

**Algorithm 3** *Online mean field variational Bayes algorithm for approximate inference in the Gaussian response linear mixed model (4).*

---

1. Perform batch-based tuning runs analogous to those described in Algorithm 2' and determine a warm-up sample size  $n_{\text{warm}}$  for which convergence is validated.
2. Set  $\mathbf{y}_{\text{warm}}$  and  $\mathbf{C}_{\text{warm}}$  to be the response vector and design matrix based on the first  $n_{\text{warm}}$  observations. Then set  $\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{C}^T \mathbf{y} \leftarrow \mathbf{C}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{C}^T \mathbf{C} \leftarrow \mathbf{C}_{\text{warm}}^T \mathbf{C}_{\text{warm}}$ ,  $n \leftarrow n_{\text{warm}}$ . Also, set  $\mu_{q(1/\sigma_{\varepsilon}^2)}$  to be the value for this quantity obtained in the batch-based tuning run with sample size  $n_{\text{warm}}$ .
3. Cycle:

read in  $\mathbf{y}_{\text{new}}$  ( $1 \times 1$ ) and  $\mathbf{c}_{\text{new}}$  ( $P \times 1$ ) ;  $n \leftarrow n + 1$   
 $\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}^T \mathbf{y} + \mathbf{y}_{\text{new}}^2$  ;  $\mathbf{C}^T \mathbf{y} \leftarrow \mathbf{C}^T \mathbf{y} + \mathbf{c}_{\text{new}} \mathbf{y}_{\text{new}}$  ;  $\mathbf{C}^T \mathbf{C} \leftarrow \mathbf{C}^T \mathbf{C} + \mathbf{c}_{\text{new}} \mathbf{c}_{\text{new}}^T$   
 $\boldsymbol{\Sigma}_{q(\beta, \mathbf{u})} \leftarrow \left[ \mu_{q(1/\sigma_{\varepsilon}^2)} \mathbf{C}^T \mathbf{C} + \text{blockdiag}\{\sigma_{\beta}^{-2} \mathbf{I}_p, \mu_{q(1/\sigma_{u1}^2)} \mathbf{I}_{K_1}, \dots, \mu_{q(1/\sigma_{ur}^2)} \mathbf{I}_{K_r}\} \right]^{-1}$   
 $\boldsymbol{\mu}_{q(\beta, \mathbf{u})} \leftarrow \mu_{q(1/\sigma_{\varepsilon}^2)} \boldsymbol{\Sigma}_{q(\beta, \mathbf{u})} \mathbf{C}^T \mathbf{y}$  ;  $\mu_{q(1/a_{\varepsilon})} \leftarrow 1 / \{\mu_{q(1/\sigma_{\varepsilon}^2)} + A_{\varepsilon}^{-2}\}$   
 $\mu_{q(1/\sigma_{\varepsilon}^2)} \leftarrow \frac{n + 1}{2 \mu_{q(1/a_{\varepsilon})} + \mathbf{y}^T \mathbf{y} - 2 \boldsymbol{\mu}_{q(\beta, \mathbf{u})}^T \mathbf{C}^T \mathbf{y} + \text{tr}[(\mathbf{C}^T \mathbf{C}) \{\boldsymbol{\Sigma}_{q(\beta, \mathbf{u})} + \boldsymbol{\mu}_{q(\beta, \mathbf{u})} \boldsymbol{\mu}_{q(\beta, \mathbf{u})}^T\}]}$   
For  $\ell = 1, \dots, r$ :  
 $\mu_{q(1/a_{u\ell})} \leftarrow 1 / \{\mu_{q(1/\sigma_{u\ell}^2)} + A_{u\ell}^{-2}\}$   
 $\mu_{q(1/\sigma_{u\ell}^2)} \leftarrow \frac{K_{\ell} + 1}{2 \mu_{q(1/a_{u\ell})} + \|\boldsymbol{\mu}_{q(\mathbf{u}_{\ell})}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\mathbf{u}_{\ell})})}$

until data no longer available or analysis terminated.

---



---

Here and throughout <sup>ind.</sup> denotes “distributed independently”.

Model (7) is the random intercept extension of simple linear regression for longitudinal data with  $(x_{ij}, y_{ij})$  denoting the  $j$ th predictor/response pair for the  $i$ th group, with  $m$  denoting the number of groups. There is no intrinsic reason to insist that the observations arrive in order with respect to the  $i, j$  subscripting. Hence  $\mathbf{c}_{\text{new}}$  will have the form

$$\mathbf{c}_{\text{new}} = \begin{bmatrix} 1 \\ \mathbf{x}_{\text{new}} \\ \mathbf{e}_{\text{new}} \end{bmatrix}.$$

Here  $\mathbf{x}_{\text{new}}$  is the new predictor measurement that partners  $y_{\text{new}}$  and  $\mathbf{e}_{\text{new}}$  is a  $m \times 1$  vector with an entry of 1 in position  $i_{\text{new}}$ , corresponding to the group that  $(x_{\text{new}}, y_{\text{new}})$  is from, and zeroes elsewhere.

Model (8) is a mixed model-based penalized spline version of the additive model

$$y_i = \beta_0 + f_s(s_i) + f_t(t_i) + \varepsilon_i, \quad 1 \leq i \leq n,$$

where the  $s_i$  and  $t_i$  are continuous predictor measurements and  $f_s$  and  $f_t$  are smooth functions. The functions  $z_k^s(\cdot)$ ,  $1 \leq k \leq K_s$ , are spline basis functions. A simple example is the truncated line basis

$$z_k^s(s) = (s - \kappa_k)_+ \tag{9}$$

where  $\kappa_1^s, \dots, \kappa_{K_s}^s$  are a set of knots within the domain of the  $s_i$  values. More sophisticated, and numerically stable, options for  $z_k(s)$  are described in, for example, Wood (2006), Welham *et al.* (2007) and Wand & Ormerod (2008). We use the last of these, known as O’Sullivan splines, in our examples. The  $z_k^t(\cdot)$ ,  $1 \leq k \leq K_t$ , are defined similarly. A key feature of the  $z_k^s(\cdot)$  and  $z_k^t(\cdot)$  is that the multiple-of-diagonal covariance matrices are appropriate under mixed model representations of penalized splines. This subtlety is explained in Section 4 of Wand & Ormerod (2008).

Online fitting of (8) involves reading in vectors of the form

$$\mathbf{c}_{\text{new}} = [1, s_{\text{new}}, t_{\text{new}}, z_1^s(s_{\text{new}}), \dots, z_{K_s}^s(s_{\text{new}}), z_1^t(t_{\text{new}}), \dots, z_{K_t}^t(t_{\text{new}})]^T$$

where  $s_{\text{new}}$  and  $t_{\text{new}}$  are the new predictor measurements that partner  $y_{\text{new}}$ . There is, however, the issue of having to set the spline basis functions in advance. For instance, if the truncated line basis (9) is used then the knots have to be set at or near the start of the algorithm. For many applications this is not a major problem. For example, if the  $s_{\text{new}}$  values correspond to age, in years, of human adults then the range of possible  $s_i$  values is easy to specify and a reasonable spline basis can be set in advance. In a similar vein, for longitudinal data, Algorithm 3 assumes that the number of groups is set in advance. If the groups correspond to the counties of a geographical entity then this should not pose a problem. If the data are from a medical study then Algorithm 3 assumes that the number of patients and their identity numbers are fixed in advance. If this is not a reasonable assumption then some adjustment is required.

Finally, we mention the possibility of speeding up the most expensive update:

$$\boldsymbol{\Sigma}_{q(\beta, \mathbf{u})} \leftarrow \left[ \mu_{q(1/\sigma_\varepsilon^2)} \mathbf{C}^T \mathbf{C} + \text{blockdiag}\{\sigma_\beta^{-2} \mathbf{I}_p, \mu_{q(1/\sigma_{u1}^2)} \mathbf{I}_{K_1}, \dots, \mu_{q(1/\sigma_{ur}^2)} \mathbf{I}_{K_r}\} \right]^{-1}. \tag{10}$$

For Model (7) the dimension of the matrix requiring inversion has dimension  $(2 + m) \times (2 + m)$ . If the number of groups is high then naïve implementation could lead to a bottleneck at (10). In the batch case it is well-known (e.g. Smith & Wand, 2008) that  $\mathbf{C}^T \mathbf{C}$  contains diagonal forms that allow  $O(m)$  computation of the right-hand side of (10). Such efficiencies are available in the online case, but require careful rearrangement of the entries of  $\mathbf{C}^T \mathbf{C}$  during the updates.

### 2.3 Extension to Unstructured Covariance Matrices for Random Effects

A *random intercepts and slopes* extension of (7) is one with the first two hierarchical levels set to

$$y_{ij} | \beta_0, U_i, \beta_1 \stackrel{\text{ind.}}{\sim} N(\beta_0 + U_i + (\beta_1 + V_i) x_{ij}, \sigma_\varepsilon^2), \quad 1 \leq i \leq m, \quad 1 \leq j \leq n_i,$$

$$\text{and } \begin{bmatrix} U_i \\ V_i \end{bmatrix} \sim N(\mathbf{0}, \boldsymbol{\Sigma}), \quad \text{where } \boldsymbol{\Sigma} \equiv \begin{bmatrix} \sigma_u^2 & \rho_{uv} \sigma_u \sigma_v \\ \rho_{uv} \sigma_u \sigma_v & \sigma_v^2 \end{bmatrix}$$

is an unstructured  $2 \times 2$  covariance matrix. The conjugate prior for  $\boldsymbol{\Sigma}$  is the Inverse Wishart distribution. However, the specification

$$\boldsymbol{\Sigma} | a_{uv1}, a_{uv2} \sim \text{Inverse-Wishart} \left( \nu + 1, 2\nu \begin{bmatrix} 1/a_{uv1} & 0 \\ 0 & 1/a_{uv2} \end{bmatrix} \right),$$

$$a_{uv1}, a_{uv2} \stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma} \left( \frac{1}{2}, 1/A_{uv} \right), \quad \nu, A_{uv} > 0$$

provides a covariance matrix extension of  $\sigma_u \sim \text{Half-Cauchy}(A_u)$ . The choice  $\nu = 2$  is particularly attractive since it imposes a Uniform( $-1, 1$ ) distribution on  $\rho_{uv}$  and Half- $t_2$  distributions on  $\sigma_u$  and  $\sigma_v$ . This is laid out in Huang & Wand (2012), including the definition of the Inverse-Wishart( $a, \mathbf{B}$ ) distribution.

Extensions to more sophisticated models, possibly having larger unstructured covariance matrices, can be done in a similar fashion.

### 2.4 Extension to Sparse Shrinkage Penalties

Model (4) involves the following Gaussian penalization on sub-vectors of  $\mathbf{u}$ :

$$\mathbf{u}_\ell | \sigma_{u\ell}^2 \sim N(\mathbf{0}, \sigma_{u\ell}^2 \mathbf{I}), \quad 1 \leq \ell \leq r. \quad (11)$$

However, many models of current-day interest, such as *wide data* (“ $p \gg n$ ”) and *wavelet* regression, require an assumption that the regression coefficients are sparse. Under such sparseness assumptions, the Gaussian priors (11) are not appropriate since they induce a relatively gentle amount of penalization that lacks the ability to annihilate regression coefficients during fitting and inference.

For simplicity of exposition we will confine discussion of the sparse shrinkage extension to the  $r = 1$  version of (4). Hence we retain

$$y | \boldsymbol{\beta}, \mathbf{u}, \sigma_\varepsilon^2 \sim N(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}, \sigma_\varepsilon^2 \mathbf{I})$$

without any sub-division of  $\mathbf{u}$ . Let  $K$  be the dimension of  $\mathbf{u}$  and consider general mutually independent prior penalization’s of the form:

$$u_k \stackrel{\text{ind.}}{\sim} p(u; \sigma_u, \boldsymbol{\theta})$$

where  $p(\cdot; \sigma_u, \boldsymbol{\theta})$  is a density function with scale parameter  $\sigma_u$  and shape parameter  $\boldsymbol{\theta}$ . Options for  $p(u; 1, \boldsymbol{\theta})$  include:

$$p(u; 1, w) = w \left\{ \frac{1}{2} \exp(-|u|) \right\} + (1 - w) \delta_0(u) \quad (\text{Laplace-Zero}),$$

$$p(u; 1) = (2\pi^3)^{-1/2} \exp(u^2/2) E_1(u^2/2) \quad (\text{Horseshoe}),$$

$$p(u; 1, \lambda) = \frac{\lambda^{2\lambda} \Gamma(\lambda + \frac{1}{2})}{\pi^{1/2}} \exp(u^2/4) D_{-2\lambda-1}(|u|) \quad (\text{Normal-Exponential-Gamma}) \quad (12)$$

$$\text{and } p(u; \lambda) = \frac{1}{2(1 + |u|/\lambda)^{\lambda+1}} \quad (\text{Generalized Double Pareto}).$$

Here  $\delta_0$  denotes the Dirac delta function with mass at zero,  $E_1$  denotes the exponential integral function of order 1 and  $D_\nu$  denotes the parabolic cylinder function of order  $\nu$  according to the definitions of Gradshteyn & Ryzhik (1994). References for the development of these sparse shrinkage priors are Johnstone & Silverman (2005) (Laplace-Zero), Carvalho, Polson & Scott (2010) (Horseshoe), Griffin & Brown (2011) (Normal-Exponential-Gamma) and Armagan, Dunson & Lee (2012) (Generalized Double Pareto).

Batch MFVB algorithms for the priors (12) have recently been derived by Wand & Ormerod (2011) (Laplace-Zero prior) and Neville, Ormerod & Wand (2012) (Horseshoe, Normal-Exponential-Gamma and Generalized Double Pareto priors).

Algorithm 4 is the online adaptation of Algorithm 4 of Wand & Ormerod (2011) for the Laplace-Zero prior model:

$$\begin{aligned}
\mathbf{y} | \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{v}, \sigma_\varepsilon^2 &\sim N(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}(\boldsymbol{\gamma} \odot \mathbf{v}), \sigma_\varepsilon^2 \mathbf{I}), \quad \mathbf{v} | \sigma_u^2, \mathbf{b} \sim N(\mathbf{0}, \sigma_u^2 \text{diag}(\mathbf{b})^{-1}), \\
\sigma_u^2 | a_u &\sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_u), \quad \sigma_\varepsilon^2 | a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/a_\varepsilon), \\
\boldsymbol{\beta} \sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{I}), \quad a_u &\sim \text{Inverse-Gamma}(\frac{1}{2}, 1/A_u^2), \quad a_\varepsilon \sim \text{Inverse-Gamma}(\frac{1}{2}, 1/A_\varepsilon^2), \\
b_k &\stackrel{\text{ind.}}{\sim} \text{Inverse-Gamma}(1, \frac{1}{2}), \quad \gamma_k | \rho \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\rho), \quad \rho \stackrel{\text{ind.}}{\sim} \text{Beta}(A_\rho, B_\rho).
\end{aligned} \tag{13}$$

Note that  $\mathbf{A} \odot \mathbf{B}$  denotes the element-wise product of matrices  $\mathbf{A}$  and  $\mathbf{B}$  having the same dimensions. Model (13) is a reproduction of (30) in Wand & Ormerod (2011) and the additional notation is explained there. Note, in particular, that the Laplace-Zero prior is handled via the introduction of auxiliary variables  $\mathbf{b}$ ,  $\boldsymbol{\gamma}$  and  $\mathbf{v}$ . Section 3.6 of Wand & Ormerod (2011) provides the necessary details. Similar online MFVB algorithms for the continuous sparse signal shrinkage priors listed in (12) follow from the batch MFVB algorithms of Neville, Ormerod & Wand (2012).

As with the spline-based semiparametric regression models described in Section 2.2, the wavelet-based models described here benefit from the *low-rank* property laid out in Section 3.1 of Wand & Ormerod (2011). This property entails that the basis functions are fixed once and for all during the warm-up period. This permits fast updating of wavelet nonparametric fits as new data arrive. A cost of this approach is that the domain of predictors needs to be specified based on the warm-up data. As explained in Section 2.2, this will often be reasonable. Of course, there is always the possibility of new predictor values landing outside domain of the basis functions, in which case some modification may be necessary.

Figure 4 illustrates online wavelet nonparametric regression for data generated to

$$x_{\text{new}} \sim \text{Uniform}(0,1), \quad y_{\text{new}} | x_{\text{new}} \sim N(f_{\text{WO}}(x_{\text{new}}), 1)$$

where  $f_{\text{WO}}$  is defined by (20) of Wand & Ormerod (2011). The warm-up sample size is  $n_{\text{warm}} = 300$ . The desired improvement in the estimate of  $f_{\text{WO}}$  as  $n$  increases is clearly apparent. Convergence to the batch MFVB estimate was found to be excellent in this case.

### 3 Binary Response Models

The binary response model we consider here takes the same form as (4) and (5), but with  $\sigma_\varepsilon$  removed and

$$\mathbf{y} | \boldsymbol{\beta}, \mathbf{u} \sim \text{Bernoulli}\{\text{logit}^{-1}(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u})\}. \tag{14}$$

Note that (14) is a convenient shorthand for the entries of  $\mathbf{y}$ , conditional on  $(\boldsymbol{\beta}, \mathbf{u})$ , being independent and with  $i$ th entry  $\text{Bernoulli}[\text{logit}^{-1}\{(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u})_i\}]$ .

Batch MFVB algorithms for approximate inference in (14), (4) and (5) start with the product restriction

$$p(\boldsymbol{\beta}, \mathbf{u}, a_{u1}, \dots, a_{ur}, \sigma_{u1}^2, \dots, \sigma_{ur}^2) \approx q(\boldsymbol{\beta}, \mathbf{u}, a_{u1}, \dots, a_{ur}) q(\sigma_{u1}^2, \dots, \sigma_{ur}^2).$$

---

**Algorithm 4** Mean field variational Bayes algorithm for the determination of the optimal parameters in  $q^*(\beta, v)$ ,  $q^*(\gamma)$ ,  $q^*(\sigma_u^2)$  and  $q^*(\sigma_\varepsilon^2)$  for the Bayesian sparse signal regression model (13).

---

1. Perform batch-based tuning runs analogous to those described in Algorithm 2' and determine a warm-up sample size  $n_{\text{warm}}$  for which convergence is validated. The batch MFVB algorithm is Algorithm 4 of Wand & Ormerod (2011).
2. Set  $\mathbf{y}_{\text{warm}}$  and  $\mathbf{C}_{\text{warm}} = [\mathbf{1} \ \mathbf{Z}_{\text{warm}}]$  to be the response vector and design matrix based on the first  $n_{\text{warm}}$  observations. Then set  $\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{Z}^T \mathbf{1} \leftarrow \mathbf{Z}_{\text{warm}}^T \mathbf{1}$ ,  $\mathbf{Z}^T \mathbf{y} \leftarrow \mathbf{Z}_{\text{warm}}^T \mathbf{y}$ ,  $\mathbf{Z}^T \mathbf{Z} \leftarrow \mathbf{Z}_{\text{warm}}^T \mathbf{Z}_{\text{warm}}$ ,  $\mathbf{C}^T \mathbf{y} \leftarrow \mathbf{C}_{\text{warm}}^T \mathbf{y}_{\text{warm}}$ ,  $\mathbf{C}^T \mathbf{C} \leftarrow \mathbf{C}_{\text{warm}}^T \mathbf{C}_{\text{warm}}$ ,  $n \leftarrow n_{\text{warm}}$ . Set  $K$  to be the number of columns in  $\mathbf{Z}_{\text{warm}}$ . Also, set  $\mu_{q(1/\sigma_\varepsilon^2)}$ ,  $\mu_{q(1/\sigma_\varepsilon^2)}$ ,  $\mu_{q(1/\sigma_u^2)}$ ,  $\mu_{q(1/a_\varepsilon)}$ ,  $\mu_{q(1/a_u)}$ ,  $\boldsymbol{\mu}_{q(b)}$ ,  $\boldsymbol{\mu}_{q(w_\gamma)}$  and  $\boldsymbol{\Omega}_{q(w_\gamma)}$  to be the values for these quantities obtained in the batch-based tuning run with sample size  $n_{\text{warm}}$ .
3. Cycle:

read in  $y_{\text{new}}$  ( $1 \times 1$ ) and  $\mathbf{z}_{\text{new}}$  ( $K \times 1$ ) ;  $n \leftarrow n + 1$  ;  $\mathbf{c}_{\text{new}} \leftarrow [1 \ \mathbf{z}_{\text{new}}]$

$$\mathbf{y}^T \mathbf{y} \leftarrow \mathbf{y}^T \mathbf{y} + y_{\text{new}}^2 ; \mathbf{Z}^T \mathbf{1} \leftarrow \mathbf{Z}^T \mathbf{1} + \mathbf{z}_{\text{new}} ; \mathbf{Z}^T \mathbf{y} \leftarrow \mathbf{Z}^T \mathbf{y} + \mathbf{z}_{\text{new}} y_{\text{new}}$$

$$\mathbf{Z}^T \mathbf{Z} \leftarrow \mathbf{Z}^T \mathbf{Z} + \mathbf{z}_{\text{new}} \mathbf{z}_{\text{new}}^T ; \mathbf{C}^T \mathbf{y} \leftarrow \mathbf{C}^T \mathbf{y} + \mathbf{c}_{\text{new}} y_{\text{new}} ; \mathbf{C}^T \mathbf{C} \leftarrow \mathbf{C}^T \mathbf{C} + \mathbf{c}_{\text{new}} \mathbf{c}_{\text{new}}^T$$

$$\boldsymbol{\Sigma}_{q(\beta, v)} \leftarrow \left( \mu_{q(1/\sigma_\varepsilon^2)} (\mathbf{C}^T \mathbf{C}) \odot \boldsymbol{\Omega}_{q(w_\gamma)} + \begin{bmatrix} \sigma_\beta^{-2} & \mathbf{0} \\ \mathbf{0} & \mu_{q(1/\sigma_u^2)} \text{diag}(\boldsymbol{\mu}_{q(b)}) \end{bmatrix} \right)^{-1}$$

$$\mu_{q(\beta, v)} \leftarrow \mu_{q(1/\sigma_\varepsilon^2)} \boldsymbol{\Sigma}_{q(\beta, v)} \text{diag}\{\boldsymbol{\mu}_{q(w_\gamma)}\} \mathbf{C}^T \mathbf{y}$$

$$\mu_{q(b)} \leftarrow [\mu_{q(1/\sigma_u^2)} \{\text{diagonal}(\boldsymbol{\Sigma}_{q(v)}) + \boldsymbol{\mu}_{q(v)}^2\}]^{-1/2}$$
  

$$\eta_{q(\gamma)} \leftarrow -\frac{1}{2} \mu_{q(1/\sigma_\varepsilon^2)} \left[ \text{diagonal}(\mathbf{Z}^T \mathbf{Z}) \odot \{\boldsymbol{\sigma}_{q(v)}^2 + \boldsymbol{\mu}_{q(v)}^2\} - 2(\mathbf{Z}^T \mathbf{y}) \odot \boldsymbol{\mu}_{q(v)} \right.$$

$$+ 2(\mathbf{Z}^T \mathbf{1}) \odot \{[\boldsymbol{\Sigma}_{q(\beta, v)}]\}_{i=1, 2 \leq j \leq K+1} + \mu_{q(\beta)} \boldsymbol{\mu}_{q(v)} \}$$

$$+ 2 \text{diagonal}\{\mathbf{Z}^T \mathbf{Z} \text{diag}(\boldsymbol{\mu}_{q(\gamma)}) \boldsymbol{\Sigma}_{q(v)}\}$$

$$- 2 \text{diagonal}(\mathbf{Z}^T \mathbf{Z}) \odot \boldsymbol{\mu}_{q(\gamma)} \odot \text{diagonal}(\boldsymbol{\Sigma}_{q(v)})$$

$$+ 2 \boldsymbol{\mu}_{q(v)} \odot \{\mathbf{Z}^T \mathbf{Z} (\boldsymbol{\mu}_{q(\gamma)} \odot \boldsymbol{\mu}_{q(v)}) - \text{diagonal}(\mathbf{Z}^T \mathbf{Z}) \odot \boldsymbol{\mu}_{q(\gamma)} \odot \boldsymbol{\mu}_{q(v)}\}$$

$$\left. + \psi(A_\rho + \mu_{q(\gamma_\bullet)}) - \psi(B_\rho + K - \mu_{q(\gamma_\bullet)}) \right]$$
  

$$\mu_{q(\gamma)} \leftarrow \frac{\exp(\eta_{q(\gamma)})}{1 + \exp(\eta_{q(\gamma)})} ; \boldsymbol{\mu}_{q(w_\gamma)} \leftarrow \begin{bmatrix} 1 \\ \boldsymbol{\mu}_{q(\gamma)} \end{bmatrix} ; \mu_{q(\gamma_\bullet)} \leftarrow \sum_{k=1}^K \mu_{q(\gamma_k)}$$

$$\boldsymbol{\Omega}_{q(w_\gamma)} \leftarrow \text{diag}\{\boldsymbol{\mu}_{q(w_\gamma)} \odot (\mathbf{1} - \boldsymbol{\mu}_{q(w_\gamma)})\} + \boldsymbol{\mu}_{q(w_\gamma)} \boldsymbol{\mu}_{q(w_\gamma)}^T$$

$$\mu_{q(1/a_\varepsilon)} \leftarrow 1/\{\mu_{q(1/\sigma_\varepsilon^2)} + A_\varepsilon^{-2}\} ; \mu_{q(1/a_u)} \leftarrow 1/\{\mu_{q(1/\sigma_u^2)} + A_u^{-2}\}$$

$$B_{q(\sigma_\varepsilon^2)} \leftarrow \mu_{q(1/a_\varepsilon)} + \frac{1}{2} \mathbf{y}^T \mathbf{y} - \left( \boldsymbol{\mu}_{q(w_\gamma)} \odot \boldsymbol{\mu}_{q(\beta, v)} \right)^T \mathbf{C}^T \mathbf{y}$$

$$+ \frac{1}{2} \text{tr} \left( \mathbf{C}^T \mathbf{C} \left[ \boldsymbol{\Omega}_{q(w_\gamma)} \odot \left\{ \boldsymbol{\Sigma}_{q(\beta, v)} + \boldsymbol{\mu}_{q(\beta, v)} \boldsymbol{\mu}_{q(\beta, v)}^T \right\} \right] \right)$$

$$B_{q(\sigma_u^2)} \leftarrow \mu_{q(1/a_u)} + \frac{1}{2} \boldsymbol{\mu}_{q(b)}^T \{\text{diagonal}(\boldsymbol{\Sigma}_{q(v)}) + \boldsymbol{\mu}_{q(v)}^2\}$$

$$\mu_{q(1/\sigma_u^2)} \leftarrow \frac{1}{2}(K+1)/B_{q(\sigma_u^2)} ; \mu_{q(1/\sigma_\varepsilon^2)} \leftarrow \frac{1}{2}(n+1)/B_{q(\sigma_\varepsilon^2)}$$

until data no longer available or analysis terminated.

---

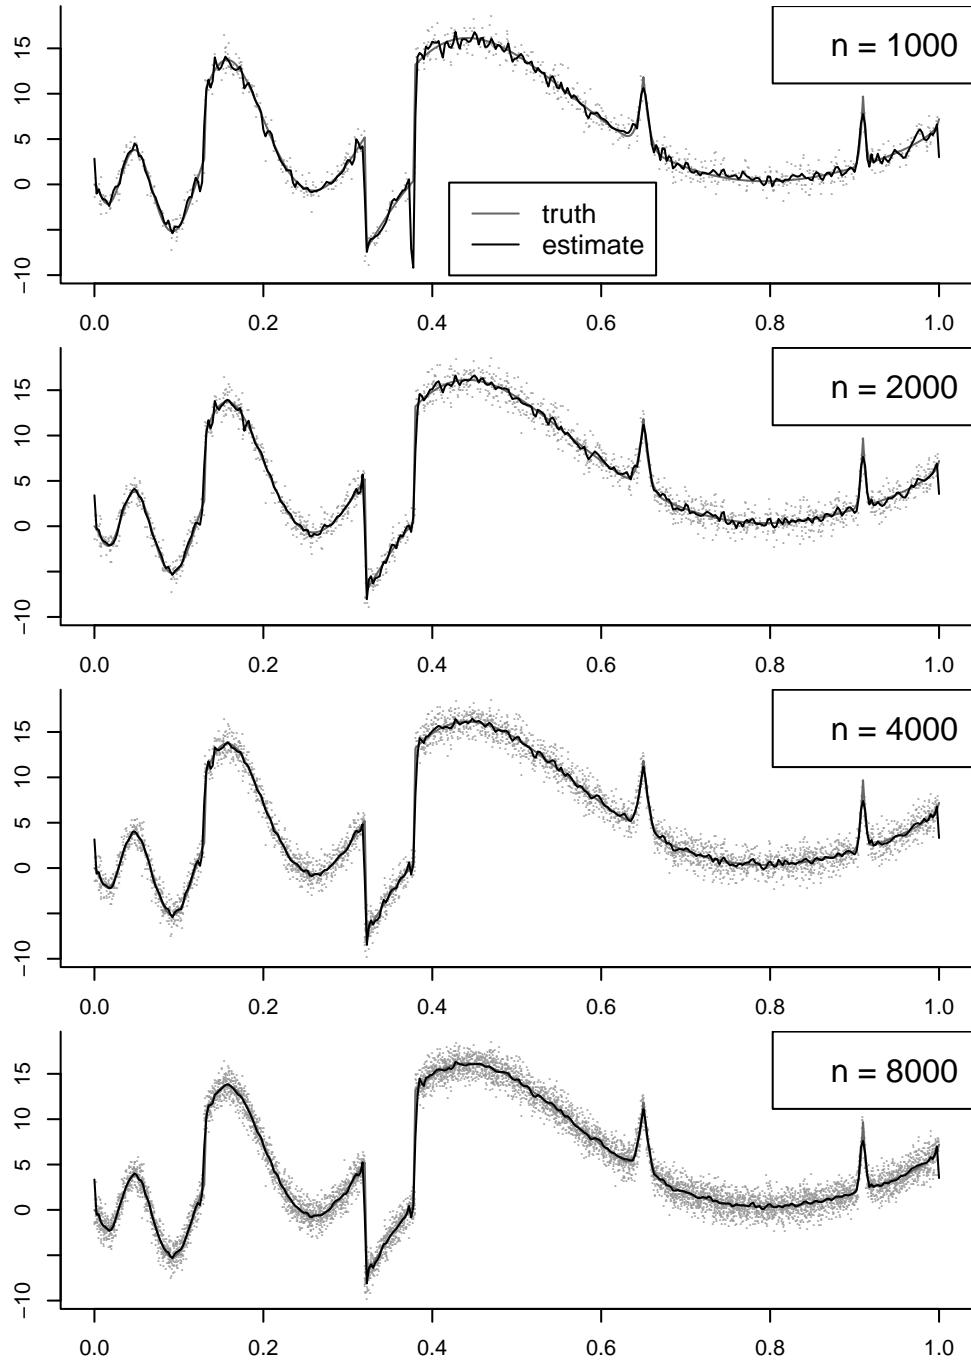


Figure 4: Examples of online MFVB wavelet fits based on Algorithm 4. The true regression curve is the function  $f_{WO}$  defined in Wand & Ormerod (2011).

The resultant updates for the  $\sigma_{u\ell}^2$  and  $a_{u\ell}$  are the same as in the Gaussian response case. The optimal  $q$ -density for  $(\beta, u)$  satisfies

$$q^*(\beta, u) \propto \exp \left\{ \mathbf{y}^T(\mathbf{X}\beta + \mathbf{Z}u) - \mathbf{1}^T \log(1 + e^{\mathbf{X}\beta + \mathbf{Z}u}) - \frac{1}{2\sigma_\beta^2} \|\beta\|^2 - \frac{1}{2} \sum_{\ell=1}^L \mu_{q(1/\sigma_{u\ell}^2)} \|u_\ell\|^2 \right\}. \quad (15)$$

However, this is a non-standard form and poses tractability problems with regards to approximate inference for  $(\beta, u)$ . A reasonable remedy is to replace (15) by a member of the following family of Multivariate Normal approximations:

$$\underline{q}^*(\beta, u) \sim N(\underline{\mu}_{q(\beta, u; \xi)}, \underline{\Sigma}_{q(\beta, u; \xi)})$$

where

$$\boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} \equiv \left[ 2\mathbf{C}^T \text{diag}\{\lambda(\boldsymbol{\xi})\} \mathbf{C} + \text{blockdiag}\{\sigma_\beta^{-2} \mathbf{I}_p, \mu_{q(1/\sigma_{u1}^2)} \mathbf{I}_{K_1}, \dots, \mu_{q(1/\sigma_{ur}^2)} \mathbf{I}_{K_r}\} \right]^{-1},$$

$\boldsymbol{\xi}$  is an  $n \times 1$  vector of positive *variational* parameters,  $\lambda(x) \equiv \tanh(x/2)/(4x)$ , and

$$\boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} \equiv \boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} \mathbf{C}^T (\mathbf{y} - \frac{1}{2} \mathbf{1})$$

with  $\mathbf{C} = [\mathbf{X} \ \mathbf{Z}]$  as before. This family of approximations is due to Jaakkola & Jordan (2000) and their genesis is given there. Section 3.1 of Ormerod & Wand (2010) explains this approximation strategy using notation similar to that used here. Jaakkola & Jordan (2000) also present an Expectation-Maximization argument that results in

$$\boldsymbol{\xi} \leftarrow \sqrt{\text{diagonal}[\mathbf{C} \{\boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})} + \boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})} \boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})}^T\} \mathbf{C}^T]}$$

being the optimal update for the  $\boldsymbol{\xi}$  vector. Algorithm 5 is the online MFVB algorithm that arises from appropriately modifying the batch MFVB algorithm for (14) with the Jaakkola & Jordan (2000) strategy.

An alternative route to an online MFVB algorithm for binary response linear mixed models involves the probit link and the Albert & Chib (1993) auxiliary variable strategy. Batch MFVB algorithms for models of this general type have been developed by Girolami & Rogers (2006) and Consonni & Marin (2007). Modification of these algorithms for the probit link version of (14) should lead to an algorithm that performs online approximate inference similar to that performed by Algorithm 5.

Figure 5 performs batch-based convergence diagnostics for a binary response nonparametric regression example. This is a special case of (14) with  $r = 1$  and  $\mathbf{Z}$  containing spline basis functions. New predictor/response pairs  $(x_{\text{new}}, y_{\text{new}})$  were generated according to

$$x_{\text{new}} \sim \text{Uniform}(0, 1), \quad y_{\text{new}} | x_{\text{new}} \sim \text{Bernoulli}(\text{logit}^{-1}(\cos(4\pi x_{\text{new}}) + 2x_{\text{new}} - 1)). \quad (16)$$

The analogues of Steps 1.-5. of Algorithm 2' were applied with an initial trial involving  $n_{\text{warm}} = 100$  and  $n_{\text{valid}} = 100$ . The Bayes estimates and 95% credible sets of the logit-transformed mean function at each of quartiles of the  $x$ -values, as well as  $\log(\sigma_u^2)$ , are shown in the upper row of Figure 5. However, they have noticeable disagreement, which indicates non-convergence of the online MFVB results to their batch counterparts and that  $n_{\text{warm}}$  should be increased. Setting  $n_{\text{warm}} = 300$  leads to the more concordant results shown in the lower row of Figure 5, indicating adequacy of this warm-up size. We have found this behaviour typical for binary response online MFVB and this simple example demonstrates the importance of batch-based tuning and convergence diagnostics.

## 4 Live Internet Demonstrations

We have launched the web-site: `realtime-semiparametric-regression.net` for displaying live real-time semiparametric regression analyses. Links on this web-site point to several examples, and we anticipate that the set of examples will grow during the next few years. At the time of this writing, the examples involve two types of real-time data: stock prices from the U.S. National Association of Securities Dealers Automated Quotations (NASDAQ) and the London Stock Exchange in the United Kingdom, and features of property rentals in Sydney, Australia.

### 4.1 Stock price data

In this set of examples, the predictor and response variable pairs correspond to pairs of stock prices. An example nonparametric regression model is

$$(\text{Microsoft stock price})_i | \boldsymbol{\beta}, \mathbf{u}, \sigma_\varepsilon^2 \stackrel{\text{ind.}}{\sim} N(\beta_0 + f((\text{Intel stock price})_i), \sigma_\varepsilon^2) \quad (17)$$

---

**Algorithm 5** *Online mean field variational Bayes algorithm for approximate inference in the binary response logistic mixed model (14).*

---

1. Perform batch-based tuning runs analogous to those described in Algorithm 2' and determine a warm-up sample size  $n_{\text{warm}}$  for which convergence is validated.
2. Set  $\mathbf{y}_{\text{warm}}$  and  $\mathbf{C}_{\text{warm}}$  to be the response vector and design matrix, and  $\boldsymbol{\xi}_{\text{warm}}$  to be the vector of variational parameters, based on the first  $n_{\text{warm}}$  observations. Then set  $\mathbf{C}^T(\mathbf{y} - \frac{1}{2}\mathbf{1}) \leftarrow \mathbf{C}_{\text{warm}}^T(\mathbf{y}_{\text{warm}} - \frac{1}{2}\mathbf{1})$ ,  $\mathbf{C}^T \text{diag}\{\lambda(\boldsymbol{\xi})\} \mathbf{C} \leftarrow \mathbf{C}_{\text{warm}}^T \text{diag}\{\lambda(\boldsymbol{\xi}_{\text{warm}})\} \mathbf{C}_{\text{warm}}$ ,  $n \leftarrow n_{\text{warm}}$ . Also, set  $\boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})}$ ,  $\boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})}$ ,  $\mu_{q(1/\sigma_{u1}^2)}, \dots, \mu_{q(1/\sigma_{ur}^2)}$  to be the values for these quantities obtained in the batch-based tuning run with sample size  $n_{\text{warm}}$ .
3. Cycle:

read in  $y_{\text{new}}$  ( $1 \times 1$ ) and  $\mathbf{c}_{\text{new}}$  ( $P \times 1$ ) ;  $n \leftarrow n + 1$

$$\begin{aligned} \xi &\leftarrow \sqrt{\mathbf{c}_{\text{new}}^T \{\boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})} + \boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})} \boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u}; \boldsymbol{\xi})}^T\} \mathbf{c}_{\text{new}}} \\ \mathbf{C}^T(\mathbf{y} - \frac{1}{2}\mathbf{1}) &\leftarrow \mathbf{C}^T(\mathbf{y} - \frac{1}{2}\mathbf{1}) + \mathbf{c}_{\text{new}} (y_{\text{new}} - \frac{1}{2}) \\ \mathbf{C}^T \text{diag}\{\lambda(\boldsymbol{\xi})\} \mathbf{C} &\leftarrow \mathbf{C}^T \text{diag}\{\lambda(\boldsymbol{\xi})\} \mathbf{C} + \lambda(\boldsymbol{\xi}) \mathbf{c}_{\text{new}} \mathbf{c}_{\text{new}}^T \\ \boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} &\leftarrow \left[ 2\mathbf{C}^T \text{diag}\{\lambda(\boldsymbol{\xi})\} \mathbf{C} + \text{blockdiag}\{\sigma_{\beta}^{-2} \mathbf{I}_p, \mu_{q(1/\sigma_{u1}^2)} \mathbf{I}_{K_1}, \dots, \mu_{q(1/\sigma_{ur}^2)} \mathbf{I}_{K_r}\} \right]^{-1} \end{aligned}$$

$$\boldsymbol{\mu}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} \leftarrow \boldsymbol{\Sigma}_{\underline{q}(\boldsymbol{\beta}, \mathbf{u})} \mathbf{C}^T(\mathbf{y} - \frac{1}{2}\mathbf{1})$$

For  $\ell = 1, \dots, r$ :

$$\begin{aligned} \mu_{q(1/a_{u\ell})} &\leftarrow 1/\{\mu_{q(1/\sigma_{u\ell}^2)} + A_{u\ell}^{-2}\} \\ \mu_{q(1/\sigma_{u\ell}^2)} &\leftarrow \frac{K_{\ell} + 1}{2\mu_{q(1/a_{u\ell})} + \|\boldsymbol{\mu}_{q(\mathbf{u}_{\ell})}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\mathbf{u}_{\ell})})} \end{aligned}$$

until data no longer available or analysis terminated.

---



---

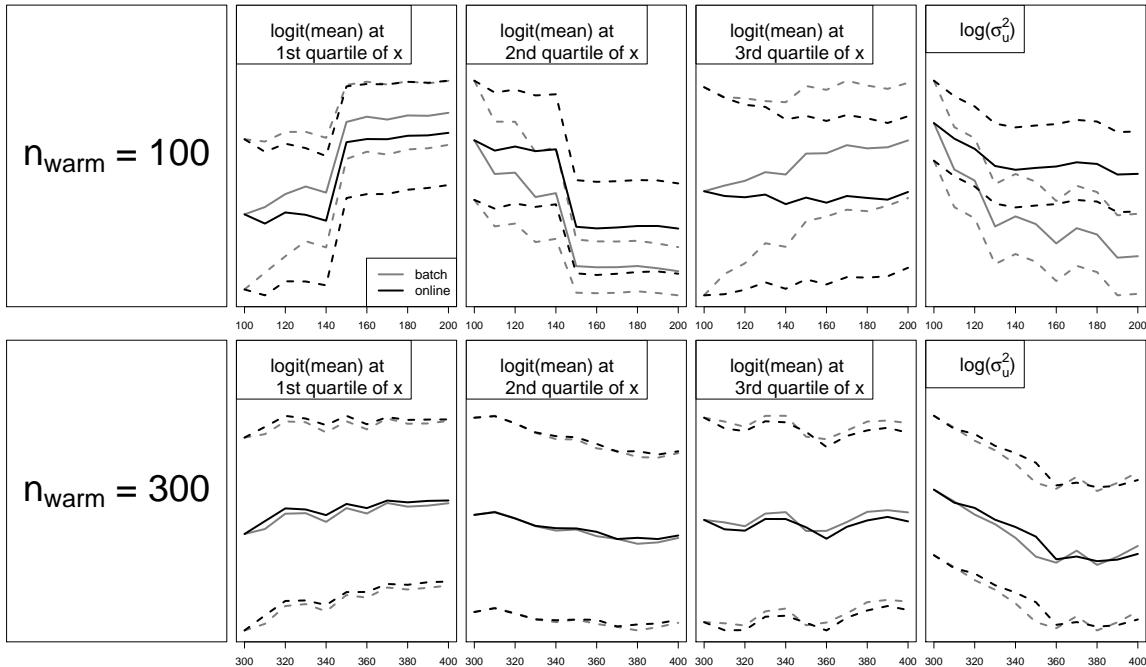


Figure 5: Convergence diagnostics for a binary response nonparametric regression example with data generated according to (16). The solid lines track the posterior means, whilst the dashed lines show corresponding 95% credible sets. First row: the horizontal axes show the sample sizes between a warm-up batch sample of size  $n_{\text{warm}} = 100$  and validation sample sizes up to  $n_{\text{valid}} = 100$  greater than  $n_{\text{warm}}$ . Second row: as for the first row, but with  $n_{\text{warm}} = 300$ .

where  $f(x) = \beta_1 x + \sum_{k=1}^K u_k z_k(x)$  is a penalized spline function as described in Section 2.2 with the same distributional structures imposed on the model parameters. In addition,  $(\text{Microsoft stock price})_i$  and  $(\text{Intel stock price})_i$  denote the  $i$ th stock price for the U.S. companies Microsoft Corporation and Intel Corporation, respectively, for the current trading day. The web-site displays fitting of (17) in real-time during the NASDAQ opening hours (9:30am to 4:00pm North American Eastern Standard Time). The R package `quantmod` (Ryan, 2012) is used to obtain the NASDAQ data from the Yahoo! Finance web-site (`finance.yahoo.com`).

A similar series of examples is set up using London Stock Exchange data during stock market opening hours (8:00 am to 4:20 pm Greenwich Mean Time). Note that Yahoo! Finance delays London Stock Exchange data by 20 minutes.

Depending on the example and the live data-set, the appropriateness of the nonparametric regression model (17) could be questionable and more sophisticated models entertained. Hence, these examples should only be viewed as simple illustrations of the concept of real-time semiparametric regression.

## 4.2 Sydney property rental data

This example involves real-time semiparametric regression analysis of data from the property rental market in Sydney, Australia. Each day, hundreds of properties come on the Sydney market and these fresh data are usually advertised on rental agency web-sites and real estate web-sites as `realestate.com.au`. This offers the possibility to perform real-time analysis and produce live and up-to-date summaries of the rental market status. An attractive approach to model such data is the special case of semiparametric regression known as geoadditive models (Kammann and Wand, 2003). Explicitly, we work with the

model

$$\begin{aligned}
\log((\text{weekly rent})_{ij}) | \boldsymbol{\beta}, U_i, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4, \mathbf{u}_5, \sigma_{\varepsilon}^2 \stackrel{\text{ind.}}{\sim} \\
N(\beta_0 + \beta_1 \text{house}_{ij} + f_2((\text{number of bedrooms})_{ij}) \\
+ f_3((\text{number of bathrooms})_{ij}) + f_4((\text{number of car spaces})_{ij}) \quad (18) \\
+ f_5(\text{longitude}_{ij}, \text{latitude}_{ij}) + U_i, \sigma_{\varepsilon}^2), \\
U_1, \dots, U_{992} | \sigma_U^2 \stackrel{\text{ind.}}{\sim} N(0, \sigma_U^2).
\end{aligned}$$

Here,  $(\text{weekly rent})_{ij}$  is the weekly rental amount in Australian dollars of the  $j$ th property for the  $i$ th real estate agency (hereafter called the  $(i, j)$ th property), and  $\text{house}_{ij}$  is an indicator of the  $(i, j)$ th property being a house, townhouse or villa (rather than an apartment). The variable  $(\text{number of bedrooms})_{ij}$  is the number of bedrooms in the  $(i, j)$ th property. Variables concerning the numbers of bathrooms and car spaces are defined similarly. The geographical location of the  $(i, j)$ th property is conveyed by the variables  $\text{longitude}_{ij}$  and  $\text{latitude}_{ij}$ . The  $U_i$ ,  $1 \leq i \leq 992$ , are random intercepts for each of the 992 agencies. The fixed effect regression coefficients  $\beta_0$ ,  $\beta_1$  and the linear contribution to  $f_2, \dots, f_5$  are stored in  $\boldsymbol{\beta}$ . Similarly, the spline basis coefficients for  $f_2, \dots, f_5$  are stored in  $\mathbf{u}_2, \dots, \mathbf{u}_5$ . The estimate of  $f_5$  is based on bivariate thin plate splines as explained in Chapter 13 of Ruppert, Wand & Carroll (2003).

The web-site for this example displays fitting of (18) in real time based on data collected since 9th May, 2012. A number of regression summaries are presented. Firstly, a geographical map is listed with processed properties as small black dots and recently (i.e. during the last hour) added ones as yellow circles. The total number of processed properties is included at the bottom right. Next, a color-coded geographical map displays the weekly rent for a two bedroom apartment with one bathroom and one car space for various geographical locations. The approximate posterior density function for  $\beta_1$  shows the impact of the property being a house or not. Regression fits and 95% credible sets for the number of bedrooms, bathrooms and car spaces for apartments are presented. Finally, a list of rental agencies with the least and most expensive properties, after correcting for all other covariates, is provided. All these regression summaries are computed in real time and the figures are updated every hour.

## Acknowledgments

This research was partially supported by Australian Research Council Discovery Project DP110100061. T. Broderick's research was supported by a National Science Foundation Graduate Research Fellowship. We are grateful to Jeff Morris and Paul Murrell for discussions related to this research.

## References

Albert, J.H. & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, **88**, 669–679.

Armagan, A., Dunson, D.B. & Lee, J. (2012). Generalized double Pareto shrinkage. *Statistica Sinica*, to appear.

Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.

Cameron, A.C. & Trivedi, P.K. (2005). *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.

Carvalho, C.M., Polson, N.G. & Scott, J.G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, **97**, 465–480.

Consonni, G. & Marin, J.-M. (2007). Mean-field variational approximate Bayesian inference for latent variable models. *Computational Statistics and Data Analysis*, **52**, 790–798.

Croissant, Y. (2011). Ecdat 0.1. Data sets for econometrics. R package, [cran.r-project.org](http://cran.r-project.org)

Devroye, L. & Wagner, T.J. (1980). On the  $L_1$  convergence of kernel estimators of regression functions with application to discrimination. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, **51**, 15–25.

Faes, C., Ormerod, J.T. & Wand, M.P. (2011). Variational Bayesian inference for parametric and nonparametric regression with missing data. *Journal of the American Statistical Association*, **106**, 959–971.

Fricker, R.D. & Chang, J.T. (2008). A spatio-temporal methodology for real-time biosurveillance. *Quality Engineering*, **20**, 465–477.

Girolami, M. & Rogers, S. (2006). Variational Bayesian multinomial probit regression. *Neural Computation*, **18**, 1790–1817.

Gradshteyn, I.S. & Ryzhik, I.M. (1994). *Tables of Integrals, Series, and Products*, 5th Edition, San Diego, California: Academic Press.

Griffin, J.E. & Brown, P.J. (2011). Bayesian hyper lassos with non-convex penalization. *Australian and New Zealand Journal of Statistics*, **53**, 423–442.

Härdle, W. (1990). *Applied Nonparametric Regression*. Cambridge: Cambridge University Press.

Huang, A. & Wand, M.P. (2012). Simple marginally noninformative prior distributions for covariance matrices. Unpublished manuscript.

Jaakkola, T.S. & Jordan, M.I. (2000). Bayesian parameter estimation via variational methods. *Statistics and Computing* **10**, 25–37.

Jank, W. & Shmueli, G. (2007). Modelling concurrency of events in on-line auctions via spatiotemporal semiparametric models. *Applied Statistics*, **56**, 1–27.

Johnstone, I.M. & Silverman, B.W. (2005). Empirical Bayes selection of wavelet thresholds. *The Annals of Statistics*, **33**, 1700–1752.

Kaimi, I. & Diggle, P.J. (2011). A hierarchical model for real-time monitoring of variation in risk of non-specific gastro-intestinal infections. *Epidemiology and Infection*, **139**, 1854–1862.

Kammann, E.E. and Wand, M.P. (2003). Geoadditive models. *Journal of the Royal Statistical Society, Series C*, **52**, 1–18.

Krzyzak, A. & Pawlak, M. (1984). Almost everywhere convergence of a recursive regression function estimate and classification. *IEEE Transactions on Information Theory*, **IT-30**, 91–93.

Michalak, S., DuBois, A., DuBois, D., Vander Wiel, S. & Hogden, J. (2012). Developing systems for real-time streaming analysis. *Journal of Computational and Graphical Statistics*, **21**, 561–580.

Neville, S.E., Ormerod, J.T. & Wand, M.P. (2012). Mean field variational Bayes for continuous sparse signal shrinkage: pitfalls and remedies. Unpublished manuscript.

Ng, S-K., McLachlan, G.J. & Lee, A.H. (2006). An incremental EM-based learning approach for on-line prediction of hospital resource utilization. *Artificial Intelligence in Medicine*, **36**, 257–267.

Ormerod, J.T. & Wand, M.P. (2010). Explaining variational approximations. *The American Statistician*, **64**, 140–153.

Ruppert, D., Wand, M.P. & Carroll, R.J. (2003). *Semiparametric Regression*. New York: Cambridge University Press.

Ruppert, D., Wand, M.P. & Carroll, R.J. (2009). Semiparametric regression during 2003–2007. *Electronic Journal of Statistics*, **3**, 1193–1256.

Ryan, J.A. (2012) quantmod 0.3. Quantitative financial modelling framework. R package, [cran.r-project.org](http://cran.r-project.org)

Smith, A.D.A.C. & Wand, M.P. (2008). Streamlined variance calculations for semiparametric mixed models. *Statistics in Medicine*, **27**, 435–448.

Tchumtchoua, S., Dunson, D.B. & Morris, J.S. (2012). Online variational Bayes inference for high-dimensional correlated data. Unpublished manuscript.

Wainwright, M.J. & Jordan, M.I. (2008). Graphical models, exponential families, and variational inference. *Foundation and Trends in Machine Learning*, **1**, 1–305.

Wand, M.P. & Jones, M.C. (1995) *Kernel Smoothing*. London: Chapman and Hall.

Wand, M.P. & Ormerod, J.T. (2008). On semiparametric regression with O’Sullivan penalized splines. *Australian and New Zealand Journal of Statistics*, **50**, 179–198.

Wand, M.P. (2009). Semiparametric regression and graphical models. *Australian and New Zealand Journal of Statistics*, **51**, 9–41.

Wand, M.P. & Ormerod, J.T. (2011). Penalized wavelets: embedding wavelets into semiparametric regression. *Electronic Journal of Statistics*, **5**, 1654–1717.

Wang, C., Paisley, J. & Blei, D.M. (2011). Online variational inference for the hierarchical Dirichlet process. *International Conference on Artificial Intelligence and Statistics*, 2011, Fort Lauderdale, Florida, USA.

Welham, S.J., Cullis, B.R., Kenward, M.G. & Thompson, R. (2007). A comparison of mixed model splines for curve fitting. *Australian and New Zealand Journal of Statistics*, **49**, 1–23.

Wolverton, C.T. & Wagner, T.J. (1969). Asymptotically optimal discriminant functions for pattern recognition. *IEEE Transactions on Information Theory*, **IT-15**, 258–265.

Wood, S.N. (2006). *Generalized Additive Models: An Introduction with R*. Boca Raton, Florida:

Chapman & Hall/CRC.

Yamato, H. (1971). Sequential estimation of a continuous probability density function and model. *Bulletin of Mathematical Statistics*, **14**, 1–12.

Zhao, Y., Staudenmayer, J., Coull, B.A. & Wand, M.P. (2006). General design Bayesian generalized linear mixed models. *Statistical Science*, **21**, 35–51.